Range-efficient consistent sampling and locality-sensitive hashing for polygons

Joachim Gudmundsson[*1] **and Rasmus Pagh**[†2]

1 **University of Sydney, Australia**
joachim.gudmundsson@sydney.edu.au
2 **IT University of Copenhagen, Denmark**
pagh@itu.dk

arXiv:1701.05290v1 [cs.CG] 19 Jan 2017

Abstract

Locality-sensitive hashing (LSH) is a fundamental technique for similarity search and similarity estimation in high-dimensional spaces. The basic idea is that similar objects should produce hash collisions with probability significantly larger than objects with low similarity. We consider LSH for objects that can be represented as point sets in either one or two dimensions. To make the point sets finite size we consider the subset of points on a grid. Directly applying LSH (e.g. min-wise hashing) to these point sets would require time proportional to the number of points. We seek to achieve time that is much lower than direct approaches. Technically, we introduce new primitives for *range-efficient* consistent sampling (of independent interest), and show how to turn such samples into LSH values. Another application of our technique is a data structure for quickly estimating the size of the intersection or union of a set of preprocessed polygons. Curiously, our consistent sampling method uses transformation to a geometric problem.

1 Introduction

Suppose that you would like to search a collection of polygons for a shape resembling a particular query polygon. Or that you have a collection of discrete probability distributions, and would like to search for a distribution that resembles a given query distribution. A framework for addressing this kind of question is *locality-sensitive hashing* (LSH), which seeks to achieve hash collisions between similar objects, while keeping the collision probability low for objects that are not very similar. Arguably the most practically important LSH method is *min-wise* hashing, which works on any type of data where similarity can be expressed in terms of *Jaccard similarity* of sets, i.e., the ratio between the size of the intersection and the size of the union of the sets. Indeed, the seminal papers of Broder et al. introducing min-wise hashing [4, 6] have more than 1000 citations. The basic idea is to map a set S to $\operatorname{argmin}_{x \in S} h(x)$, which for a strong enough random hash function h gives collision probability equal to (or very close to) the Jaccard similarity (see e.g. [5] for a discussion of sufficient requirements on h).

If we represent discrete probability distributions by histograms there is a one-to-one relationship between the Jaccard similarity of two histograms and the statistical distance between the corresponding distributions, so a search for close distributions in terms of Jaccard similarity will translate into a search for distributions that are close in statistical distance, see Figure 1.

To make min-wise hashing well-defined on infinite point sets in the plane we may shift to an approximation by considering only those points contained in a finite grid of points. However, for a good approximation these sets must be very large, which means that computing a hash value $h(x)$ for each point $x \in S$, in order to do min-wise hashing, is not attractive.

* Supported under Australian Research Council's Discovery Projects funding scheme (project number DP150101134).
† Supported by the European Research Council under the European Union's 7th Framework Programme (FP7/2007-2013) / ERC grant agreement no. 614331.

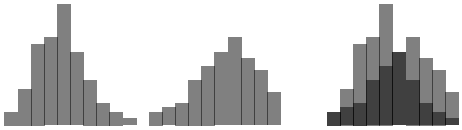

Figure 1 Two probability distributions represented as histograms (left). Their statistical distance can be computed from the Jaccard similarity of the point sets (right).

1.1 Our results

We consider efficient locality-sensitive hashing for objects that can be represented as point sets in either one or two dimensions, and whose similarity is measured as the Jaccard similarity of these point sets. The model of computation considered is a Word RAM with word size at least $\log p$, where p is a prime number. We use integers in $U = \{0, \ldots, p-1\}$ (or equivalently elements in the field $\mathcal{F}_p$ of size p) to represent coordinates of points on the grid. Our first result concerns histograms with n values in U:

▸ **Theorem 1.** *For every constant $\varepsilon > 0$ and every integer N it is possible to choose an explicit hash function $H : U^n \to D$ that has constant description size, can be evaluated in time $O(n \log p)$, and for which $\Pr[H(\mathbf{x}) = H(\mathbf{y})] \in [J - \varepsilon; J + \varepsilon]$, where $J = \frac{\sum_i \min(x_i, y_i)}{\sum_i \max(x_i, y_i)}$ is the weighted Jaccard similarity of vectors $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{y} = (y_1, \ldots, y_n)$ of weight $\sum_i x_i = \sum y_i = N$.*

Our construction gives an explicit alternative to existing results on weighted min-wise hashing (see [14, 16]) whose analysis relies on hash functions that are fully random and cannot be described in small space. It was previously shown that a form of priority sampling based on 2-independence can be used to estimate Jaccard similarity of histograms [21], but similarity estimation is less general than locality-sensitive hashing methods such as weighted min-wise hashing.

We proceed to show the generality of our technique by presenting an LSH method for geometric objects. We will use approximation to achieve high performance even for "hard" shapes, and adopt the so-called *fuzzy model*, proposed by Aryan and Mount [1]. Consider a polygon P to be a fuzzy polygon, and assume that points that are "close" to the boundary of the polygon (relative to the polygon's diameter) may or may not be included in the polygon. That is, given a polygon P and real value $0 < \phi \leq 1$, define the outer range $P^+ = P^+(w)$ to be the locus of points whose distance from a point interior to P is at most $w = \phi \cdot d(P)$, where $d(P)$ is the diameter of P. The inner range $P^- = P^-(w)$ of P is defined symmetrically.

Using the fuzzy model a *valid answer* to the Jaccard similarity of two polygons P_1 and P_2 with respect to ϕ is any value $\frac{X_\cap}{X_\cup}$ such that $A(P_1^- \cap P_2^-) \leq X_\cap \leq A(P_1^+ \cap P_2^+)$ and $A(P_1^- \cup P_2^-) \leq X_\cup \leq A(P_1^+ \cup P_2^+)$, where $A(\cdot)$ denotes the area of the region. To simplify the statement of the theorem we say that a polygon is *α-dense* in a rectangle I if for some constant $\alpha > 0$ its area is at least a fraction α of the area of I.

▸ **Theorem 2.** *For every choice of constants $\varepsilon > 0$, $\phi > 0$, $\alpha > 0$ and square $I \subseteq \mathbf{R}^2$ it is possible to choose an explicit random hash function H whose description size is constant, that can be evaluated in time $O(t \log p)$, where t is the time to test if a given point lies inside a polygon, and with the following guarantee on collision probability: Let $P_1, P_2 \subseteq I$ be polygons such that P_1^+ and P_2^+ are α-dense in I. Then $\Pr[H(P_1) = H(P_2)] \in [J - \varepsilon; J + \varepsilon]$, where J is some valid Jaccard similarity of P_1 and P_2 in the fuzzy model with parameter ϕ.*

In Section 5 we present further applications of our technique and show how a small summary can be constructed for a set $\mathcal{P}$ of polygons such that for any subset $\mathcal{Q}$ of $\mathcal{P}$, an estimate of the area of $\cap\mathcal{Q}$ and $\cup\mathcal{Q}$ can be computed efficiently in the fuzzy model with respect to ϕ.

Techniques.

Our main technical contribution lies in methods for *range-efficient min-wise hashing* in one and two dimensions, efficiently implementing min-wise hashing for intervals and rectangles. More specifically, we consider intervals in U and rectangles in $U \times U$. The new technique can be related to earlier methods for *sampling* items with small hash values in one or more dimensions [19, 22]. (In fact, en route we obtain new hash-based sampling algorithms with improved speed, which may be of independent interest.) However, using [19, 22] to sample a single item is not likely to yield a good locality-sensitive hash function. The reason is that the hash functions used in these methods are taken from simple, 2-independent families and, as explained by Thorup [21], min-wise hashing using 2-independence does not in general yield collision probability that is close to (or even a function of) the Jaccard similarity. Instead we use a 2-phase approach: First produce a sample of k elements having the smallest hash values, and then perform standard min-wise hashing on a carefully selected *subset* of the sample using a *different* hash function.

We can combine and filter the samples to handle a variety of point sets that are not intervals or rectangles. To create a sample for a subset of a rectangle we can generate a (larger) sample of the rectangle, and then filter away those sample points that are not in the subset. This is efficient if the subset is suitably dense in the rectangle (which we may ensure by working in the fuzzy model). To create a sample from the union of two sets, simply take the union of the samples. Theorems 1 and 2 are obtained in this way, and it would be possible to instantiate many other applications.

At the heart of our range-efficient sampling algorithms for one and two dimensions lies a reduction to the problem of finding an integer point (or integer points) in a given interval with small vertical distance to a given line. Such a point can effectively be found by traversing the integer convex hull of the line. Using a result of Charrier and Buzer [9] this can be done in logarithmic time. Thus, geometry shows up in an unexpected way in the solution.

1.2 Comparison with related work

We are not aware of previous work dealing with range-efficient locality-sensitive hashing. The most closely related work is on range-efficient *consistent* (or *coordinated*) sampling, which is a technique for constructing summaries and sketches of large data sets. The technique comes in two flavors: *bottom-k* (or min-wise) sampling, which fixes the sample size, and *consistent sampling* (or *sub-sampling*), which fixes the sampling probability. In both cases the idea is to choose as a sample those elements from a set $S \subseteq U$ that have small hash values under a random hash function $h : U \to [0; 1]$. If the sample size is fixed and some hash values are identical an arbitrary tiebreaking rule can be used, e.g., selecting the minimum element. To make argmin uniquely defined, which is convenient, we take $\text{argmin}_{x \in I}\, h(x)$ to be the smallest value $y \in I$ for which $h(y) = \min_{x \in I} h(x)$. To denote the set of the k elements having the smallest hash values (with ties broken in the same way) we use the notation argmin_k. We focus on settings in which U is large and it is not feasible to store a table of all hash values, or other data structures of similar size.

In one dimension.

Pavan and Tirthapura [19] consider the 2-independent family of linear hash functions in the field of size p, i.e., functions of the form $h(x) = (ax + b) \bmod p$. They show how to efficiently find hash values $h(x)$ below a given threshold Δ, where x is restricted to an interval I. Their algorithm uses time $O(\log p + k)$, where k is the number of elements $x \in I$ with $h(x) \leq \Delta$. It seems that using this in connection with doubling search might lead to an algorithm finding the minimum hash value in time $O(\log^2 p)$. In this paper we show how to improve the time complexity:

▸ **Lemma 3.** *Let $h(x) = (ax + b) \bmod p$, where p is prime and $0 \le a, b < p$. Given $i_2 > i_1 > 0$ consider the interval $I = \{i_1, \dots, i_2\}$. It is possible to compute* $\operatorname{argmin}_{x \in I} h(x)$ *(the* min-hash *of I) in time $O(\log |I|)$.*

We will argue in Section 2 that Lemma 3 can be applied repeatedly to subintervals to output the k smallest hash values (and corresponding inputs) in time $O(k \log |I|)$. The possibility of choosing $a = 0$ is included for mathematical convenience (to ensure 2-independence), though in most applications it will be better to choose $a > 0$ (which in addition makes argmin uniquely defined without a tiebreaking rule).

In more than one dimension.

Tirthapura and Woodruff [22] consider another class of 2-independent functions, namely linear transformations on vectors over the field $\mathcal{F}_2$. Integers naturally correspond to such vectors, and for a *dyadic* interval I containing all integers that share a certain prefix, the problem of finding elements in I that map to zero is equivalent to solving a linear system of equations. Since an arbitrary interval can be split into a logarithmic number of dyadic intervals they are able to compute all the integers that map to zero in polylogarithmic time. The sampling probability can be chosen as an arbitrary integer power of two. This method generalizes to rectangles in dimension $d \ge 2$.

In this paper we instead consider 2-independent functions of the form

$$(x, y) \mapsto (ax + by + c) \bmod p \ .$$

We do not know of a method for efficiently computing a min-hash over a rectangle for such functions, but we are able to efficiently implement consistent sampling with sampling probability $1/p$.

▸ **Lemma 4.** *Let $h(x, y) = (ax + by + c) \bmod p$, where p is prime and $0 \le a, b, c < p$. Given $i_1 < i_2$ and $j_1 < j_2$ consider $I = \{i_1, \dots, i_2\} \times \{j_1, \dots, j_2\}$. It is possible to compute $I' = \{(x, y) \in I \mid h(x, y) = 0\}$ in time $O((|I'| + 1) \log(\min(i_2 - i_1, j_2 - j_1)))$.*

For random a, b, c the expected size of the sample I' is $|I|/p$, and because of 2-independence the distribution of $|I'|$ is concentrated around this value. Compared to the method of [22] ours is faster, but has the disadvantage that the sampling probability cannot be chosen freely. However, as we will see this restriction is not a real limitation to our applications to locality sensitive hashing and size estimation.

From consistent sampling to LSH.

Our technique for transforming a consistent sample to an LSH value is of independent interest. Thorup [21] shows that min-wise hashing using 2-independence does not in general yield collision probability that is close to (or even a function of) the Jaccard similarity. On the positive side he shows that bottom-k *samples* of two sets made using a 2-independent hash function can be used to estimate the Jaccard similarity J of the sets with arbitrarily good precision. However, this does not yield a locality-sensitive hash function with collision probability (close to) J, and obvious approaches such as min-wise hashing applied to the samples fails to have the right collision probability. Instead, we use consistent sampling (using a 2-independent family) followed by a stronger hash function for which min-wise hashing has the desired collision probability up to an additive error ε. This transformation yields the first LSH family for Jaccard similarity (with proven guarantees on collision probability) where the function can be:

- evaluated in $O(n + \text{poly}(1/\varepsilon))$ time on a set of size n, and
- described and computed in a constant number of machine words (independent of n).

Previous such functions have used either time per element that grows as ε approaches zero [15], or required description space that is a root of n (see [11]).

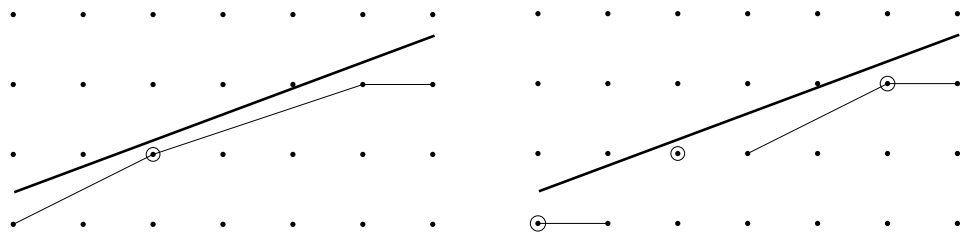

Figure 2 Illustration of reduction to integer convex hull

1.3 Preliminaries

We will make extensive use of 2-independence:

▸ **Definition 5.** A family of hash functions $\mathcal{H}$ mapping U to U is called *2-independent* if $\forall x_1, x_2, a_1, a_2 \in U$ with $x_1 \neq x_2$ and $h \in \mathcal{H}$ chosen uniformly we have

$$\Pr[h(x_1) = a_1 \wedge h(x_2) = a_2] = 1/|U|^2 \ .$$

It will be convenient to use the notation $x \pm \Delta$ for a number in the interval $[x - \Delta; x + \Delta]$.

Carter and Wegman [7] showed that the family $\mathcal{H}_1 = \{x \mapsto (ax + b) \bmod p \mid a, b \in U\}$ is 2-independent on the set $U = \{0, \ldots, p-1\}$ when p is a prime. Finally, we make use of ε-minwise independent families:

▸ **Definition 6.** A family of hash functions $\mathcal{H}$ mapping U to $\mathbb{N}$ is called *ε-minwise independent* if for every set $S \subset U$, every $y \in S$, and random $h \in \mathcal{H}$: $\Pr[h(y) = \min h(S)] = (1 \pm \varepsilon)/|S|$.

Indyk [15] showed that an efficient ε-minwise independent family mapping to a range of size $O(p/\varepsilon)$ can be constructed by using an $O(\log(1/\varepsilon))$-independent family of functions (e.g. polynomial hash functions). Dahlgaard and Thorup [11] showed that the evaluation time can be made constant, independent of ε, by using space $|U|^{\Omega(1)}$. If we only care about sets of size up to some number $\hat{n}$, this space usage can be improved to $(\hat{n}/\varepsilon)^{\Omega(1)}$.

2 Range-efficient bottom-k sampling in one dimension

The aim of this section is to show Lemma 3 and how it can be used to efficiently compute consistent as well as bottom-k samples. Together with the general transformation presented in Section 4 this will lead to Theorem 1.

Without loss of generality suppose $0 < i_1 < i_2 < p$, consider $I = \{i_1, \ldots, i_2\} \subseteq U$, and let $h \in \mathcal{H}_1$. To show Lemma 3 we must prove that $\operatorname{argmin}_{x \in I} h(x)$ can be computed in time $O(\log |I|)$. In case $a = 0$ this is trivial (just output i_1), so we focus on the case $a > 0$. We will show how the problem can be reduced to the problem of finding the integer point at the smallest (vertical) distance below the line segment

$$\ell = \{(x, (ax + b)/p) \mid x \in [i_1; i_2]\}. \tag{1}$$

To see this observe that for $x \in \mathbb{N}$ we have vertical distance $(ax + b)/p - \lfloor (ax + b)/p \rfloor$ between the line and the nearest integer point. Using the equality

$$(ax + b)/p - \lfloor (ax + b)/p \rfloor = ((ax + b) \bmod p)/p$$

we see that minimizing $(ax + b \bmod p)$ is equivalent to minimizing $(ax + b)/p - \lfloor (ax + b)/p \rfloor$, as claimed. Therefore it suffices to search for the point $(x, y) \in D = I \times \mathbb{N}$ below ℓ that is closest

to ℓ. Since ℓ is a line, the point (x, y) must lie on the convex hull $CH(\ell)$ of the set of points in D that lie below ℓ, referred to as the "integer convex hull", see Figure 2. Clearly, the closest point will always be on the upper part of the hull, denoted $CH_L(\ell)$. Zolotykh [24] showed that $CH(\ell)$ consists of $O(\log(i_2 - i_1))$ line segments. To find a point on the integer convex hull with the smallest vertical distance to ℓ we will use a result by Charrier and Buzer [9].

▸ **Theorem 7.** *(Charrier and Buzer [9]) Given a line segment ℓ, the upper integer convex hull $CH_L(\ell)$ can be computed in $O(\log(i_2 - i_1))$ time, where i_1 and i_2 are the x-coordinates of the end points of ℓ.*

Charrier and Buzer initially assume that ℓ passes through the origin. However, they note (Section 7 in [9]) that this requirement is not needed. Thus, using their result on the line ℓ defined in (1) we obtain Lemma 3.

We now discuss how to use Lemma 3 to output the k smallest hash values (and corresponding inputs, i.e., the bottom-k sample) in time $O(k \log p)$. First compute $CH_L(\ell)$ and find the point $(x_1, y_1) \in CH_L(\ell)$ with the smallest vertical distance to ℓ. Next, split the problem into two subintervals; one for the part of ℓ in the x-interval $[i_1, x_1 - 1]$ and one for the part of ℓ in the x-interval $[x_1 + 1, i_2]$. Using a heap to find the integer point with smallest vertical distance in the intervals considered, we can repeat this process until k points have been found. To compute a consistent sample rather than the bottom-k sample we simply stop the procedure whenever we see an element with a hash value larger than the threshold.

▸ **Corollary 8.** *Let $h(x) = (ax + b) \bmod p$, where p is prime and $0 \le a, b < p$. Given $0 \le i_1 < i_2 < p$ consider $I = \{i_1, \ldots, i_2\}$. It is possible to compute the bottom-k sample (or the consistent sample of expected size k) from the interval I with respect to h in (expected) time $O(k \log |I|)$.*

3 Rectangle-efficient consistent sampling

The aim and structure of this section are similar to those of Section 2, but now addressing the case where we want to do hashing-based sampling in a rectangle $I = \{i_1 \ldots, i_2\} \times \{j_1 \ldots, j_2\}$. Specifically, we prove Lemma 4 and show how one can use it to perform consistent sampling. This will be used in Section 4 to prove Theorem 2 and in Section 5 to construct an efficient data structure for estimating the size of intersections and unions of polygons. Assume without loss of generality that $0 \le i_1 < i_2 < p$, $0 \le j_1 < j_2 < p$ and $i_2 - i_1 \le j_2 - j_1$. Consider the 2-independent family $\mathcal{H}_2 = \{(x, y) \mapsto (ax + by + c) \bmod p \mid a, b, c \in U\}$ and choose $h \in \mathcal{H}_2$. To prove Lemma 4 we have to argue that

$$I' = \{(x, y) \in I \mid h(x, y) = 0\} \tag{2}$$

can be computed in time $O((|I'| + 1) \log(i_2 - i_1))$. Similar to the previous section we will show how the problem can be reduced to the problem of finding all integer points below a line segment ℓ with a small vertical distance to ℓ.

To find all $(x, y) \in I$ for which $h(x, y) = (ax + by + c) \bmod p = 0$, as a first step we translate the function h such that we can consider input $y \in [0, y_2']$. Specifically, we replace h with $h' : (x, y) \mapsto (ax + by + c') \bmod p$, where $c' = c + bj_1$, and consider inputs with $(x, y) \in [i_1, i_2] \times [0, j_2']$, $j_2' = j_2 - j_1$. This is equivalent to the original task since $h(x, y) = h'(x, y - j_1)$. Next note that for $x \in [i_1, i_2]$ and $y \in [0, j_2']$:

$$(ax + by + c') \bmod p = 0 \quad \Leftrightarrow \quad y \equiv (-b^{-1}ax - b^{-1}c') \bmod p,$$

To simplify the expression set $q = -b^{-1}a$ and $s = -b^{-1}c'$. Then we have a zero hash value when $y = (qx + s) \bmod p = (qx + s) - kp$ for some positive integer k. Dividing by p and substituting

$\alpha = q/p$ and $\beta = s/p$ we get:

$$\frac{y}{p} = \alpha x + \beta - k,$$

where $x \in [i_1, i_2]$ and $y/p \in [0, j_2'/p]$. Now we can express the original problem as finding all $(x,k) \in [i_1, i_2] \times \mathbb{N}$ such that $\alpha x + \beta - k \in [0, j_2'/p]$. Consider the line segment

$$\ell' = \{(x, \alpha x + \beta) \mid x \in [i_1; i_2]\}. \tag{3}$$

An integer point (x', y') below ℓ' with $x' \in [i_1, i_2]$ and vertical distance at most j_2'/p to ℓ' corresponds to a point (x', y') such that $h'(x', y') = 0$.

To find all the points (x', y') that fulfill the restrictions we can apply the same technique as in Section 2. That is, compute the integer convex hull $CH_L(\ell')$ using the algorithm by Charrier and Buzer [9]. One difference from the setting of Section 2 is that we are interested in *all* integer points close to ℓ', but $CH_L(\ell')$ is guaranteed only to include one such point if it exists. This is handled by recursing on subintervals in which no points have been reported until we find an interval where the integer convex hull does not contain a point close to ℓ'. Recall that the time to output the integer convex hull is $O(\log(i_2 - i_1))$ by the result of Zolotykh [24], so the cost per point reported is logarithmic. This concludes the proof of Lemma 4.

3.1 Concentration bound

▸ **Definition 9.** An (ε, δ)-estimator for a quantity μ is a randomized procedure that, given parameters $0 < \varepsilon < 1$ and $0 < \delta < 1$, computes an estimate X of μ such that $\Pr[|X - \mu| > \varepsilon\mu] < \delta$.

For some $\alpha > 0$ consider an arbitrary set $S \subseteq I$, and the sample $S' = S \cap I'$ where I' is defined in (2). We now show that $p|S'|$ is concentrated around its expectation $|S|$ when p is not too large.

▸ **Lemma 10.** *For $\varepsilon > 0$, $p|S'|$ is an $(\varepsilon, p/(\varepsilon^2\mu))$-estimator for $\mu = |S|$.*

Proof. The proof is a standard application of the second moment bound for 2-independent indicator variables. For each point $p \in S$ let X_p be the indicator variable that equals 1 if $p \in I'$ and 0 otherwise. Clearly we have $|S'| = X$ where $X = \sum_{p \in S} X_i$, so $\mathbf{E}[pX] = p\sum_{p \in S} \mathbf{E}[X_i] = \mu$. By definition of I' the variables are 2-independent, and so $\mathbf{Var}(pX) = p^2\mathbf{Var}[X] \le p^2\mathbf{E}[X] = p\mu$. Now Chebychev's inequality implies $\Pr[|pX - \mu| > \varepsilon\mu] < \mathbf{Var}(pX)/(\varepsilon\mu)^2 \le p/(\varepsilon^2\mu)$. ◂

To get an (ε, δ)-estimator we thus need $p \le \delta\varepsilon^2\mu$. The expected time for computing I' in Lemma 4 is upper bounded by $O(\mathbf{E}[|I'| + 1]\log p)$ which is $O((|I|/p + 1)\log p)$. If we choose $p = \Omega(\delta\varepsilon^2|S|)$, to get an (ε, δ)-estimator, and let $\alpha = |S|/|I|$ be the fraction of points of I that are also in S, then the expected time simplifies to $O(\log(p)/(\alpha\delta\varepsilon^2))$. That is, the bound independent of the size of S, has logarithmic dependence on p, and linear dependence on $1/\alpha$, $1/\delta$, and $1/\varepsilon$.

4 From consistent sampling to locality-sensitive hashing

We now present a general transformation of methods for 2-independent consistent sampling to locality-sensitive hashing for Jaccard similarity. Together with the consistent sampling methods in Sections 2 and 3 this will yield Theorems 1 and 2.

Thorup [21] observed that min-wise hashing based on a 2-independent family does not in general give collision probability that is close to (or a function of) Jaccard similarity. In fact, he observes bias for a 2-independent family of hash functions based on multiplication, similar to the ones used in this paper. Thus we take a different route: First produce a consistent sample using 2-independence, and then apply min-wise hashing *to the sample* using a stronger hash function. The expected time per element is constant if we make sure that the sample has expected constant size.

Let constants $\varepsilon > 0$ and $\alpha > 0$ be given. For a point set $S \subseteq I$ with $|S| \geq \alpha|I|$ we produce a 2-independent sample $I' \cap S$ with sampling probability $1/p^*$, where $p^* = \Theta(\varepsilon^3\alpha|I|)$. This is possible even when p^* is restricted to being a prime number (as in the previous sections), because there exists a prime p_i in every interval $\{2^i, \ldots, 2^{i+1} - 1\}$, $i = 1, 2, 3, \ldots$. Now select f at random from an $\varepsilon/4$-minwise independent family and define the hash value

$$H^*(S) = \operatorname*{argmin}_{x \in I' \cap S} f(x) \ . \tag{4}$$

▸ **Lemma 11.** *For $S, T \subseteq I$ with $|S|, |T| \geq \alpha|I|$ we have $\Pr[H^*(S) = H^*(T)] = \frac{|S \cap T|}{|S \cup T|} \pm \varepsilon$, where the probability is over the choice of I' and f.*

Proof. Let $S' = S \cap I'$ and $T' = T \cap I'$ denote the samples. Consider the Jaccard similarity of the samples

$$J' = \frac{|S' \cap T'|}{|S' \cup T'|} = \frac{|S'| + |T'| - |S' \cup T'|}{|S' \cup T'|} \ .$$

Conditioned on a fixed I', the collision probability of $H^*(S)$ is $J' \pm \varepsilon/4$ by the choice of f. Thus it suffices to show that J' differs from J by at most $\varepsilon/2$ with probability at least $1 - \varepsilon/4$

By Lemma 10, $p \cdot |S' \cup T'|$ is an $(\varepsilon/8, \varepsilon/12)$-estimator for $|S \cup T|$ if we increase p by a suitable constant factor. Similarly, $p \cdot |S'|$ is an $(\varepsilon/8, \varepsilon/12)$-estimator for $|S|$ and $p \cdot |T'|$ is an $(\varepsilon/8, \varepsilon/12)$-estimator for $|T|$. The probability that all estimators are good is at least $1 - \varepsilon/4$, and in that case

$$J - \varepsilon/2 < \frac{|S \cap T| - (3\varepsilon/8)|S \cup T|}{|S \cup T| + (\varepsilon/8)|S \cup T|} \leq J' \leq \frac{|S \cap T| + (3\varepsilon/8)|S \cup T|}{|S \cup T| - (\varepsilon/8)|S \cup T|} < J + \varepsilon/2$$

as desired. ◂

We have not specified how to choose f. The most obvious choice is to use an $O(\log(1/\varepsilon))$-independent hash function [15]. Another appealing choice is twisted tabulation hashing [11] that yields constant evaluation time, independent of ε. The expected size of $S \cap I'$ is bounded by a function of ε and α. This means that we can combine twisted tabulation with an injective universe reduction step to reduce the domain of twisted tabulation to a (large) constant depending on ε and α.

Proof of Theorem 1. Consider a vector $\mathbf{x} = (x_1, \ldots, x_n) \in U^n$. We follow the folklore approach [14] of conceptually mapping each vector $\mathbf{x}$ to a set $P_{\mathbf{x}}$, such that the Jaccard similarity of $P_{\mathbf{x}}$ and $P_{\mathbf{y}}$ exactly equals the weighted Jaccard similarity of $\mathbf{x}$ and $\mathbf{y}$. In particular, it is easy to verify that this is the case if we let $P_{\mathbf{x}} = \{(i, j) \mid i = 1, \ldots, n;\ j = 1, \ldots, x_i\}$. Note that $P_{\mathbf{x}}$ and $P_{\mathbf{y}}$ both have size N. We will use the following class of hash functions from $U \times U$ to U:

$$H_2 = \{(x, y) \mapsto (ax + by + c) \bmod p \mid a, b, c \in U\} \ . \tag{5}$$

The 2-independence of H_2 follows from the arguments of Carter and Wegman [7]. A proof is included in Appendix A for completeness. When restricted to points of the form $(i, \cdot)$ for a fixed i, each function $h \in \mathcal{H}_2$ has a form suitable for Corollary 8 in Section 2. This means we can find the minimum for $P_{\mathbf{x}}$ restricted to a given column i in time $O(\log x_i)$. Using a heap to keep track of the smallest hash value from each column of $P_{\mathbf{x}}$ not (yet) reported in the sample, we can output all elements of $P_{\mathbf{x}}$ with a hash value smaller than any given threshold τ in time $O(\log p)$ per element. The threshold τ is chosen to match the desired sampling probability p^*.

Lemma 11 then says that we get the desired collision probability up to an additive error of ε. The expected time to hash is $O(n \log p)$ (to populate the priority queue) plus $O(\log p)$ times the expected number of samples. The expected number of samples $|S|/p$ is constant for every constant $\varepsilon > 0$, which gives the desired time bound in expectation.

It is possible to turn the expected bound into a worst case bound by stopping the computation if the running time exceeds $1/\delta$ times the expectation, which happens with probability at most δ. If we simply output a constant in this case the collision probability changes by at most δ. Adjusting constants, this extra collision probability is inconsequential. ◂

Proof of Theorem 2. The proof is similar to the proof of Theorem 1 but with some added geometric observations. Let P_1 and P_2 be two polygons contained in I. As mentioned in the introduction, a *valid answer* to the Jaccard similarity of polygons P_1 and P_2 with respect to ϕ is any value $\frac{X_\cap}{X_\cup}$ such that $A(P_1^-(w_1) \cap P_2^-(w_2)) \le X_\cap \le A(P_1^+(w_1) \cap P_2^+(w_2))$ and $A(P_1^-(w_1) \cap P_2^-(w_2)) \le X_\cup \le A(P_1^+(w_1) \cap P_2^+(w_2))$, where $w_i = \phi \cdot d(P_i)$ for $i \in \{1,2\}$.

We now switch to considering the restrictions of $P_1^+(w_1/2)$ and $P_2^+(w_1/2)$ to a p-by-p grid of points whose enclosing rectangle contains I. The grid points are identified in the natural way with integer coordinates in $[p] \times [p]$. We choose p such that the number of points inside I is $\Theta(p/\alpha)$ times the desired number of samples required for Lemma 11 to hold.

Let $L^+ = [i_1, i_2] \times [j_1, j_2]$ be the minimum bounding box of $I \cap P_1^+(w_1/2)$ and $I \cap P_2^+(w_2/2)$. Since P_1 and P_2 are α-dense in I there exists a constant $\alpha > 0$ such that $|P_1^+(w_1/2)|, |P_2^+(w_2/2)| > \alpha \cdot |L^+|$. The consistent sampling will be made on $P_i^+(w_i/2)$, $i \in \{1,2\}$. The reason for this is that

$$|P_1^+(w_1/2) \cap P_2^+(w_2/2)| / |P_1^+(w_1/2) \cup P_2^+(w_2/2)|$$

is a valid answer to the Jaccard similarity of P_1 and P_2 in the fuzzy model with respect to ϕ, which follows immediately from the below two inequalities that are proven in Lemma 13 (Section 5):

$$A(P_1 \cup P_2) \le |(P_1^+(w_1/2) \cup P_2^+(w_2/2)) \cap I| \le A(P_1^+(w_1/2) \cup P_2^+(w_2/2))$$

and

$$A(P_1 \cap P_2) \le |(P_1^+(w_1/2) \cap P_2^+(w_2/2)) \cap I| \le A(P_1^+(w_1/2) \cap P_2^+(w_2/2)).$$

Lemma 11 gives us the desired collision probability up to an additive error of ε. The expected time to hash is $O(\log p)$ plus $O(t \log p)$ times the expected number of samples, where t is the time to test if a given grid point lies inside a polygon. The expected number of samples $|L^+|/p$ is constant for any constants ε, α, and ϕ, which gives the desired time bound in expectation.

5 Estimating union and intersection of polygons

In this section we consider the question: Given a set $\mathcal{P} = \{P_1, \ldots, P_n\}$ of n preprocessed polygons in the plane, how efficiently can we compute the area of the union or the intersection of a given subset $\mathcal{Q} \subseteq \mathcal{P}$?

Computing the union of a set of geometric objects is a well-studied problem in computational geometry. One example is the Klee's Measure Problem (KMP). Given n axis-parallel boxes in the d-dimensional space, the problem asks for the measure of their union. In 1977, Victor Klee [17] showed that it can be solved in $O(n \log n)$ time for $d = 1$. This was generalized to $d > 1$ dimensions by Bentley [2] in the same year, and later improved by van Leeuwen and Wood [23], Overmars and Yap [18] and, Chan [8]. In 2010, Bringmann and Friedrich [3] gave an $O(\frac{dn}{\epsilon^2})$ Monte Carlo $(1+\epsilon)$-approximation algorithm for the problem. They also proved $\#P$-hardness for axis-parallel boxes.

A related question is the computation of the area of the intersection of n polygons. For general boxes the problem can be shown to be $\#P$-hard [3]. Bringmann and Friedrich showed that there cannot be a (deterministic or randomized) multiplicative $(2^{d^{1-\epsilon}})$-approximation algorithm in general, unless NP=BPP. They therefore gave an additive ϵ-approximation for a large class of geometric bodies, with a running time of $O(\frac{nd}{\epsilon^2})$ assuming that the following three queries can be

approximately answered efficiently: point inside body, volume of body and sample point within a body.

In this section we will approach the problem slightly differently. The approach we suggest is to produce a small summary of the set $\mathcal{P}$, such that given any subset $\mathcal{Q}$ of $\mathcal{P}$ the union and intersection of $\mathcal{Q}$ can be estimated efficiently. Unfortunately, the lower bound arguments by Bringmann and Friedrich defeat any reasonable hope of achieving efficient performance for arbitrary polygons. To get around the lower bounds we again adopt the approximation model proposed by Aryan and Mount [1] (stated in Section 1.1), which has been used extensively in the literature [10, 12, 13].

Similar to the approach by Bringmann and Friedrich [3] we will also use sampling of the polygons to estimate the size of the union and intersection. However, compared to earlier attempts, the main advantage of our approach is that we generate the sample points (a summary of the input) in a preprocessing step and after that we may discard the polygons. Union and intersection queries are answered using only the summary. Also, we do not impose any restrictions on the input polygons. The drawbacks are that we only consider the case when $d = 2$ and the approximation model [1] we use is somewhat more "forgiving" than previously used models.

For each polygon P_i in $\mathcal{P}$, $1 \leq i \leq n$, let $w_i = \phi \cdot d(P_i)$, where $d(P_i)$ is the diameter of P_i and $0 \leq \phi \leq 1$ is a given constant. Let $\mathcal{Q}$ be a the input to a union or intersection query, that is, $\mathcal{Q}$ is a subset of $\mathcal{P}$. To simplify the notations we will write $\cup\mathcal{Q}^+(w) = \cup_{P_i \in \mathcal{Q}} P_i(w_i)$ and $\cup\mathcal{Q}^-(w) = \cup_{P_i \in \mathcal{Q}} P_i(w_i)$. Define $\cap\mathcal{Q}^+(w)$ and $\cap\mathcal{Q}^-(w)$ symmetrically.

Following the above discussion, given a legal answer to a set intersection query $X = \cap\mathcal{Q}$ is any X' such that

$$\cap\mathcal{Q}^-(w) \subseteq X' \subseteq \cap\mathcal{Q}^+(w)$$

and for a union query $X = \cup\mathcal{Q}$ a legal answer is any X' such that

$$\cup\mathcal{Q}^-(w) \subseteq X' \subseteq \cup\mathcal{Q}^+(w).$$

It is immediate from the above definitions that for any polygon P and any $w \geq \sqrt{2}$ we have:

$$P^-(w) \subset P \subset P^+(w).$$

We will use the number of integer coordinates, denoted $|P|$, within a polygon P to estimate the area of the polygon, denoted $A(P)$.

▸ **Lemma 12.** *For a polygon P having integer coordinates we have $A(P) \leq |P|$.*

Proof. We use Pick's theorem [20]. Let $i(P)$ be the number of integer coordinates in the interior of P and let $b(P)$ be the number of integer coordinates on the boundary of P. Pick's theorem states:

$$A(P) = i(P) + \frac{b(P)}{2} - (h + 1),$$

where h is the number of holes in P. Since $A(P) < i(P) + b(P) = |P|$ the lemma follows. ◂

To make the queries more efficient we will not estimate the number of integer coordinates in the intersection/union X of a query, instead we will estimate an approximation of $|X|$. For this we need the following lemma.

▸ **Lemma 13.** *For any polygon P and any $w \geq 2\sqrt{2}$, we have:*

$$A(P) \leq A(P^+(w/2)) \leq |P^+(w/2)| \leq A(P^+(w)).$$

Proof. The first inequality is immediate and the second inequality follows from Lemma 12. For the third inequality we note that any point in the plane within distance $\sqrt{2}$ from an integer coordinate within $P^+(w/2)$ will lie within $P^+(w)$, hence $|P^+(w/2)| \leq A(P^+(w))$. This completes the proof of the lemma. ◂

As an immediate consequence of Lemma 13 we can use the consistent samples in $P_i^+(w_i/2)$, $1 \le i \le n$, for our estimates of the intersection and union, provided that $w_i \ge 2\sqrt{2}$. It remains to show how a summary of $\mathcal{P}$ can be computed and how the summary can be used to answer union and intersection queries.

5.1 Constructing a summary

For a given query $\mathcal{Q}$ containing $k \le n$ polygons, let $P_{\min} = \text{argmin}_{P_i \in \mathcal{Q}} |P_i^+(w_i/2)|$, $P_{\max} = \text{argmax}_{P_i \in \mathcal{Q}} |P_i^+(w_i)/2|$ and let $d_{\min} = \text{argmin}_{P_i \in \mathcal{Q}} d(P_i)$. If $P_i = P_{\max}$ and $P_j = P_{\min}$ then we will write $P_{\max}^+(w) = P_i^+(w_i)$ and $P_{\min}^+(w) = P_j^+(w_i)$, respectively. Before stating the construction of the summary and the query algorithms we note the following two lemmas.

▸ **Lemma 14.**

$$|P_{\max}^+(w/2)| \le |\cup \mathcal{Q}^+(w/2)| \le k \cdot |P_{\max}^+(w/2)|.$$

▸ **Lemma 15.** *If* $\cap Q^-(w) \ne \varnothing$ *and* $\phi \cdot d_{\min} > 2\sqrt{2}$ *then*

$$\frac{\phi^2}{2} \cdot |P_{\min}^+(w/2)| \le |\cap \mathcal{Q}^+(w/2)| \le |P_{\min}^+(w/2)|.$$

Proof. Since $\cap\mathcal{Q}^+(w/2) \subseteq P_{\min}^+(w/2)$ the second inequality is immediate. For the first inequality we first observe that

$$A(P_{\min}^+(w/2)) \le \frac{\pi}{2}((1+\phi) \cdot d_{\min})^2. \tag{6}$$

To see this let a and b be two points on the boundary of $P_{\min}^+$ with largest inter-point distance among all points in $P_{\min}^+(w/2)$. The distance between a and b is at most $(1+\phi) \cdot d_{\min}$. Note that $P_{\min}^+(w/2)$ is enclosed in the intersection of the two disks of radius $|ab|$ centered at a and b, hence, (6) follows.

Next consider a point p in the non-empty set $\cap\mathcal{Q}^-(w)$. By definition p must lie in $P_i^-(w)$ for every $P_i \in \mathcal{Q}$. As a result the ball of radius $\frac{3}{2}\phi \cdot d_{\min}$ and center at p must lie in $\cap\mathcal{Q}$, thus

$$A(\cap\mathcal{Q}^+(w/2)) \ge \pi(\frac{3}{2} \cdot \phi \cdot d_{\min})^2. \tag{7}$$

Using the same argument as in the proof of Lemma 13 we can show:

$$|P_{\min}^+(w/2)| \le A(P_{\min}^+(w/2+\sqrt{2}) \le \frac{\pi}{2}((1+\phi) \cdot d_{\min})^2 < \pi((1+\phi) \cdot d_{\min})^2, \tag{8}$$

where the second inequality follows from (6).

Now we are ready to prove the first inequality of the statement of the lemma.

$$\begin{aligned}
|\cap \mathcal{Q}^+(w/2)| &\ge A(\cap\mathcal{Q}^+(w/2)) && \text{[from Lemma 13]}\\
&\ge \pi(\frac{3}{2} \cdot \phi \cdot d_{\min})^2 && \text{[from (7)]}\\
&> \frac{\phi^2}{2} \cdot \pi((1+\phi) \cdot d_{\min})^2 &&\\
&> \frac{\phi^2}{2} \cdot |P_{\min}^+(w/2)| && \text{[from (8)]}
\end{aligned}$$

This completes the proof of the lemma. ◂

We will use the rectangle-efficient consistent sampling technique described in Section 3 to generate a summary of $\mathcal{P}$ that can be used to estimate the area of $\cap\mathcal{Q}$ or $\cup\mathcal{Q}$, where $\mathcal{Q}$ is a given subset of $\mathcal{P}$.

The idea of the construction algorithm for the summary is simple. Let $X = \cap\mathcal{Q}^+(w/2)$ or $X = \cup\mathcal{Q}^+(w/2)$ depending on the query and, assume that $\phi \cdot d_{\min} > 2\sqrt{2}$. In a preprocessing step construct a summary of $\mathcal{P}$, denoted $\mathcal{S}$. The summary $\mathcal{S}$ will contain consistent samples for a number of different sampling rates. To answer a query, pick a minimum sampling rate $1/p$ that guarantees that the expected number of consistent samples in X is small but sufficient to guarantee an (ε, δ)-estimate of $|X|$. If X contains enough unique consistent samples then the algorithm reports an estimate of X, otherwise it iteratively increases the sampling rate with a constant factor until X contains sufficiently many unique consistent samples. From the analysis in Section 3.1 we know that an (ε, δ)-estimator of X requires the sampling rate to be approximately $1/(\delta \cdot \varepsilon^2 \cdot |X|)$.

From Lemmas 14 and 15 we have that the smallest area that will ever be considered in a query $\mathcal{Q}$ has size at least $f_{\min} = \frac{\phi^2}{2}|P^+_{\min}(w/2)|$ and the largest area is at most $f_{\max} = n \cdot |P^+_{\max}(w/2)|$. To get an (ε, δ)-estimate of $|X|$ at least $1/\delta^2\varepsilon$ unique consistent samples are required to lie within X. Now we can state the construction algorithm of the summary (c_1 and c_2 are appropriate constants).

1. Set $f := c_1 \cdot f_{\max}$
2. while $f > c_2 \cdot f_{\min}$ do
   a. $\ell := \lfloor \log f \rfloor$
   b. $p[\ell] :=$ a prime number in $[2^{\ell-1}, 2^\ell)$
   c. For every $P_i \in \mathcal{P}$ do
      i. Compute the consistent samples within the rectilinear bounding box of $P_i^+(w_i/2)$ using Lemma 4. If the number of consistent samples within $P_i^+(w_i/2)$ is greater than zero then store the sample points in $\mathcal{S}[s, j]$.
   d. $f := f/2$
3. Return $p[\cdot]$ and $S[\cdot, \cdot]$

As output from the above algorithm we get two data structures:

- $p[\ell]$: Returns a prime number between $[2^{\ell-1}, 2^\ell]$.
- $\mathcal{S}[P_i, \ell]$: Returns the set of consistent samples within $P_i^+(w_i/2)$ for which $(ax+by+c) \bmod p[j] = 0$, if the set is non-empty. Otherwise it returns False.

Complexity

Consider the total number of consistent samples generated for a polygon P_i. The number of consistent samples is expected to increase with a factor of two in each iteration of the algorithm, that is, the expected total number of consistent samples form an exponentially growing geometric series which sums to $O(\frac{1}{\phi^2\delta^2\varepsilon} \cdot \frac{|P_i|}{|P_{\min}|})$. Summing up over all the polygons, the total number of consistent samples is bounded by $O(\frac{n}{\phi^2\delta^2\varepsilon} \cdot \frac{|P_{\max}|}{|P_{\min}|})$, which is also the expected size of the summary.

For the time complexity we first note that the above procedure can be implemented such that iterations where no consistent samples are expected to be generated are omitted without consideration. Since at least a fraction of $\phi/2$ of all consistent samples in the minimal bounding box of $P_i^+(w_i/2)$ is expected to lie within $P_i^+(w_i/2)$ (can be shown using a similar argument as in the proof of Lemma 15) the total number of generated consistent samples is expected to be at most a factor of $2/\phi$ greater than the number of consistent samples in the summary. Each consistent sample requires at most $O(\log|P_i|)$ time to generate, according to Theorem 4. If we assume that testing if a consistent sample lies inside a polygon can be done in time t then the expected time to build a summary of P is $O(\frac{n}{\phi^2\delta^2\varepsilon} \cdot \frac{|P_{\max}|}{|P_{\min}|} \cdot (t + \log|P_{\max}|))$.

5.2 Union queries

A union query is a subset $\mathcal{Q}$ of $\mathcal{P}$ containing $k \leq n$ polygons. Let $f^{\cup}_{\min} = |P^{+}_{\max}(w/2)|$, let $f^{\cup}_{\max} = k \cdot |P^{+}_{\max}(w/2)|$ and, let $X = \cup Q^{+}(w/2)$. The search for a good size consistent sample starts by selecting a $j \in \mathbb{Z}$ such that $2^{j-1} \leq c \cdot \delta\varepsilon^2 \cdot f^{\cup}_{max} \leq 2^j$, for some constant c. The expected total number of consistent samples, $\sum_{P_i \in P} \mathcal{S}[P_i, p[j]]$, is $O(1/\delta\varepsilon^2)$. This is the initial sampling rate tested by the algorithm. If the number of unique consistent samples in X is at least $\frac{1}{\delta\varepsilon^2}$ then the algorithm returns the estimate $|X|_S \cdot p[j]$, where $|X|_S$ is the number of unique consistent samples in X. Otherwise, the algorithm increases the sampling probability iteratively with a factor of approximately two, by decreasing j by one, until the number of unique consistent samples in X is at least $\frac{1}{\delta\varepsilon^2}$. At that point the algorithm returns the estimate $|X|_S \cdot p[j]$.

In each iteration the number of unique consistent samples is $O(1/\delta\varepsilon^2)$, which implies that the expected total number of consistent samples considered in one iteration is $O(k/\delta\varepsilon^2)$. As noted above the number of consistent samples increases by roughly a factor of two in each iteration, thus the total number of consistent samples considered can be described by an exponentially growing geometric function which has an upper bound of $O(k/\delta\varepsilon^2)$. The size of the union of consistent samples can be computed in expected linear time with respect to the total number of consistent samples, thus $O(k/\delta\varepsilon^2)$ time in total.

5.3 Intersection queries

An intersection query is handled similarly, but instead of fetching the consistent samples in all the polygons, we only fetch the consistent samples in $P^{+}_{\min}(w/2)$ since each unique consistent sample in $\cap\mathcal{Q}^{+}(w/2)$ must also be in $P^{+}_{\min}(w/2)$. Then, for each consistent sample $s \in \mathcal{S}[P^{+}_{\min}(w/2), \ell]$, check if s is in $\mathcal{S}[P^{+}_{i}(w_i/2), \ell]$ for all $P_i \in \mathcal{Q}$. The membership query can be answered in constant expected time using hashing, thus the time required for one iteration is k times the number of consistent samples in $P^{+}_{\min}(w/2)$.

Let $f^{\cap}_{\min} = \frac{\phi^2}{2}|P^{+}_{\min}(w/2)|$, let $f^{\cap}_{\max} = |P^{+}_{\min}(w/2)|$ and, let $X = \cap Q^{+}(w/2)$. Again we start by selecting a $j \in \mathbb{Z}$ such that $2^{j-1} \leq c \cdot \gamma\delta\varepsilon^2 \cdot f^{\cap}_{max} \leq 2^j$, for some constant c.

The expected total number of consistent samples in $P^{+}_{i}(w/2)$ is $O(1/\delta\varepsilon^2)$. This is the initial sampling rate tested by the algorithm. If the number of consistent samples in X is at least $\frac{1}{\delta\varepsilon^2}$ then the algorithm returns the number of consistent samples in X times $p[j]$. Otherwise, the algorithm increases the sampling probability iteratively with a factor of approximately two, by decreasing j by one, until the number of unique consistent samples in X is at least $\frac{1}{\delta\varepsilon^2}$.

Since $f^{\cap}_{\min} = \frac{\phi^2}{4}|P^{+}_{\min}(w/2)|$, the expected total number of consistent samples considered is bounded by $\frac{2}{\phi^2} \cdot \frac{1}{\delta\varepsilon^2}$). To summarize the query time adds up to $O(\frac{k}{\phi^2\delta\varepsilon^2})$).

We can now summarize the results in this section:

▸ **Theorem 16.** . *Given a set $\mathcal{P} = \{P_1, \ldots, P_n\}$ of polygons and three constants $\varepsilon, \delta > 0$ and $0 < \phi \leq 1$. If $\phi \cdot d(P_i) \geq 2\sqrt{2}$ for all $P_i \in \mathcal{P}$ then, in the fuzzy model with respect to ϕ, there exists a summary of size $O(\frac{n}{\phi^2\delta^2\varepsilon} \cdot \frac{|P_{\max}|}{|P_{\min}|} \cdot (t + \log|P_{\max}|))$ such that for any subset $\mathcal{Q}$ of $\mathcal{P}$ containing $k \leq n$ polygons an (ε, δ)-estimate of $\cup Q$ can be computed in $O(k/\delta\varepsilon^2)$ expected time and an (ε, δ)-estimate of $\cap Q$ can be computed in $O(\frac{k}{\phi^2\delta\varepsilon^2})$ expected time.*

6 Conclusion and open problems

We have investigated efficient methods for consistent sampling and locality-sensitive hashing of 2-dimensional point sets. Though the methods are simple, it is not clear if they are as useful in practice as, say, minwise hashing. In addition to practicality, some theoretical questions remain, for example whether the additive constant ε in our theorems can be avoided. Further, our measure of similarity among point sets is by no means the only possible one — it would be interesting,

for example, to consider notions of similarity that are invariant under rotations and translations of the point set.

References

1 S. Arya and D. M. Mount. Approximate range searching. *Computational Geometry – Theory and Applications*, 17:135–152, 2000.

2 J. L. Bentley. Algorithms for Klee's rectangle problems. Unpublished note, Computer Science Department, Carnegie Mellon University, 1977.

3 K. Bringmann and T. Friedrich. Approximating the volume of unions and intersections of high-dimensional geometric objects. *Computational Geometry – Theory and Applications*, 43(6-7):601–610, 2010.

4 A. Z. Broder. On the resemblance and containment of documents. In *Proceedings of the International Conference on Compression and Complexity of Sequences (SEQUENCES)*, pages 21–29. IEEE, 1997.

5 A. Z. Broder, M. Charikar, A. M. Frieze, and M. Mitzenmacher. Min-wise independent permutations. *Journal of Computer and System Sciences*, 60(3):630–659, 2000.

6 A. Z. Broder, S. C. Glassman, M. S. Manasse, and G. Zweig. Syntactic clustering of the web. *Computer Networks and ISDN Systems*, 29(8):1157–1166, 1997.

7 J. L. Carter and M. N. Wegman. Universal classes of hash functions. In *Proceedings of the 9th ACM Symposium on Theory of Computing (STOC)*, pages 106–112. ACM, 1977.

8 T. M. Chan. Klee's measure problem made easy. In *Proceedings of the 54th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 410–419, 2013.

9 E. Charrier and L. Buzer. Approximating a real number by a rational number with a limited denominator: A geometric approach. *Discrete Applied Mathematics*, 157:3473–3484, 2009.

10 D. Z. Chen, M. H. M. Smid, and B. Xu. Geometric algorithms for density-based data clustering. *International Journal of Computational Geometry and Applications*, 15(3):239–260, 2005.

11 S. Dahlgaard and M. Thorup. Approximately minwise independence with twisted tabulation. In *Proceedings of the 14th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT)*, pages 134–145. Springer International Publishing, 2014.

12 T. K. Dang. Solving approximate similarity queries. *Computer Systems Science and Engineering*, 22(1-2):71–89, 2007.

13 S. A. Friedler and D. M. Mount. *Spatio-temporal Range Searching over Compressed Kinetic Sensor Data*, pages 386–397. Springer Berlin Heidelberg, 2010.

14 B. Haeupler, M. Manasse, and K. Talwar. Consistent weighted sampling made fast, small, and easy. *arXiv:1410.4266*, 2014.

15 P. Indyk. A small approximately min-wise independent family of hash functions. *Journal of Algorithms*, 38(1):84–90, 2001.

16 S. Ioffe. Improved consistent sampling, weighted minhash and L1 sketching. In *Proceedings of the 10th IEEE International Conference on Data Mining (ICDM)*, pages 246–255, 2010.

17 V. Klee. Can the measure of $\cup[a_i, b_i]$ be computed in less than $o(n \log n)$ steps? *American Mathematical Monthly*, 84:284–285, 1977.

18 M. H. Overmars and C.-K. Yap. New upper bounds in Klee's measure problem. *SIAM Journal on Computing*, 20(6):1034–1045, 1991.

19 A. Pavan and S. Tirthapura. Range-efficient counting of distinct elements in a massive data stream. *SIAM Journal on Computing*, 37(2):359–379, 2007.

20 G. Pick. Geometrisches zur Zahlenlehre. *Sitzungsberichte des deutschen naturwissenschaftlich-medicinischen Vereines für Böhmen "Lotos", Prag*, 19:311–319, 1889.

21 M. Thorup. Bottom-k and priority sampling, set similarity and subset sums with minimal independence. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 371–380. ACM, 2013.

22 S. Tirthapura and D. Woodruff. Rectangle-efficient aggregation in spatial data streams. In *Proceedings of the 31st Symposium on Principles of Database Systems (PODS)*, pages 283–294. ACM, 2012.

23 J. van Leeuwen and D. Wood. The measure problem for rectangular ranges in d-space. *Journal of Algorithms*, 2(3):282–300, 1981.

24 N. Y. Zolotykh. On the number of vertices in integer linear programming problems. Technical report, University of Nizhni Novograd, 2000.

A 2-independence of $\mathcal{H}_2$

▸ **Lemma 17.** *The family $\mathcal{H}_2$ defined in (5) is 2-independent.*

Proof. Let us check that our hash function is in fact 2-independent. Let $x_1, x_2, y_1, y_2 \in U$ and $t_1, t_2 \in \mathcal{F}_p$ s.t. $q_1 = (x_1, y_1) \neq q_2 = (x_2, y_2)$. Let $q^{-1} \in \mathcal{F}_p$ be the unique multiplicative inverse of q. Note that this is guaranteed to exist if and only if q is non-zero. What is the probability that $h(q_1) = t_1$ and $h(q_2) = t_2$? We have $(ax_1 + by_1 + c) \bmod p = t_1$ and $(ax_2 + by_2 + c) \bmod p = t_2$. Since $q_1 \neq q_2$ we may assume without loss of generality that $x_1 \neq x_2$. Now fix b. We get:

$$\begin{pmatrix} 1 & x_1 \\ 1 & x_2 \end{pmatrix} \begin{pmatrix} a \\ c \end{pmatrix} = \begin{pmatrix} t_1 - by_1 \\ t_2 - by_2 \end{pmatrix}$$

For every b there exists exactly one pair a, c such that the above equality holds. Since a and c are drawn uniformly and independently from $\mathcal{F}_p$, this probability is $1/p^2$. ◂ ◂